



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/980,920	04/11/2002	Michael R. Krause	10002166-2	3383

22879 7590 08/10/2007
HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

EXAMINER

NEURAUTER, GEORGE C

ART UNIT	PAPER NUMBER
----------	--------------

2143

MAIL DATE	DELIVERY MODE
-----------	---------------

08/10/2007

PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES

Application Number: 09/980,920
Filing Date: April 11, 2002
Appellant(s): KRAUSE ET AL.

MAILED

AUG 10 2007

Technology Center 2100

Patrick G. Billig, Reg. No. 38,080
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 27 June 2007
appealing from the Office action mailed 27 February 2007.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

(7) Claims Appendix

Art Unit: 2143

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

5,991,797	FUTRAL et al	11-1999
6,360,220	FORIN	3-2002
6,647,423	REGNIER et al	11-2003

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless -

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims 1-2, 6-7, 9-13, 16-17, 19-20, 23-29, and 32-33 are rejected under 35 U.S.C. 102(e) as being anticipated by US Patent 5 991 797 to Futral et al.

Regarding claim 1, Futral discloses a method of managing memory in a distributed computer system, the method comprising:

Art Unit: 2143

binding ("register") a remote key ("memory handle") to a first address representing a contiguous memory address range accessible by a first consumer process stored in a first memory at a first host processor endnode including a first processor and the first memory; (column 5, lines 18-42, specifically lines 24-34)

sending the bound remote key and first address from the first host processor endnode to a second host processor endnode on a communication fabric ("SAN fabric") via a first networking interface controller (NIC) in the first host processor endnode and a second NIC in the second host processor endnode, wherein the second host processor endnode includes a second processor and a second memory; (column 5, lines 18-42, specifically lines 32-34) and

performing a remote direct memory access operation from the second host processor endnode with a second consumer process stored in the second memory to access the contiguous memory address range including sending the bound remote key and the first address from the second host processor endnode to the first host processor endnode on the communication fabric via the second NIC and the first NIC. (column 5, lines 18-42, specifically lines 31-42; column 7, lines 5-17)

Art Unit: 2143

Regarding claim 2, Futral discloses the method of claim 1 further comprising controlling local memory access protection in the first host processor endnode with a virtual memory manager in an operating system kernel process stored in the first memory. (column 4, lines 51-65; column 5, lines 6-23)

Regarding claim 6, Futral discloses the method of claim 1 wherein the first address is an effective address ("virtual address") pointing to an address space in the first memory accessible by the first consumer process. (column 3, lines 50-58; column 5, lines 5-17 and 39-56)

Regarding claim 7, Futral discloses the method of claim 6 wherein the effective address points to a virtual address space. (column 3, lines 50-58; column 5, line 43-56)

Regarding claim 9, Futral discloses the method of claim 1 wherein the first consumer process is a user process. ("application program"; column 1, lines 8-11)

Regarding claim 10, Futral discloses the method of claim 1 wherein the first consumer process is a kernel process. ("virtual interface"; column 3, lines 15-58; column 4, lines 51-65)

Regarding claim 11, Futral discloses the method of claim 1 wherein the first address is a virtual address accessible by the first consumer process which is a consumer kernel process.

Art Unit: 2143

(column 3, lines 50-58; column 4, lines 51-65; column 5, lines 5-17 and 39-56)

Regarding claim 12, Futral discloses the method of claim 1 wherein the binding includes associating the first address to the remote key with a consumer process employing a bind remote key verb. (column 5, lines 18-42, specifically lines 29-31)

Regarding claim 13, Futral discloses the method of claim 1 further comprising obtaining at least one remote key with a consumer process employing an allocate remote key verb. (column 5, lines 18-42, specifically lines 31-32)

Regarding claim 16, Futral discloses the method of claim 1 wherein the remote key cannot be used to protect more than one memory region at a given instant. (column 5, lines 39-42)

Regarding claim 17, Futral discloses the method of claim 1 further comprising reusing the remote key after the remote direct memory access operation from the second host processor endnode is completed. (column 5, lines 23-42)

Claims 19-20, 23-29, and 32-33 are also rejected since these claims recite a distributed computer system that contains substantially the same limitations as recited in claims 1-2, 6-7, 9-13, and 16-17 respectively.

Claim Rejections - 35 USC § 103

Art Unit: 2143

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

The factual inquiries set forth in *Graham v. John Deere Co.*, 383 U.S. 1, 148 USPQ 459 (1966), that are applied for establishing a background for determining obviousness under 35 U.S.C. 103(a) are summarized as follows:

1. Determining the scope and contents of the prior art.
2. Ascertaining the differences between the prior art and the claims at issue.
3. Resolving the level of ordinary skill in the pertinent art.
4. Considering objective evidence present in the application indicating obviousness or nonobviousness.

This application currently names joint inventors. In considering patentability of the claims under 35 U.S.C. 103(a), the examiner presumes that the subject matter of the various claims was commonly owned at the time any inventions covered therein were made absent any evidence to the contrary.

Applicant is advised of the obligation under 37 CFR 1.56 to point out the inventor and invention dates of each claim that was not commonly owned at the time a later invention was made in

Art Unit: 2143

order for the examiner to consider the applicability of 35 U.S.C. 103(c) and potential 35 U.S.C. 102(e), (f) or (g) prior art under 35 U.S.C. 103(a).

Claims 3-5, 8, 21-22, and 25 are rejected under 35 U.S.C. 103(a) as being unpatentable over Futral et al. in view of US Patent 6 647 423 to Regnier et al.

Regarding claim 3, Futral discloses the method of claim 1.

Futral does not expressly disclose the method further comprising comparing the bound remote key and the corresponding first address supplied by the second host processor endnode to the bound remote key and corresponding first address in the first host processor endnode, however, Futral does disclose using the bound remote key and the corresponding first address for validation purposes (column 5, lines 32-42; column 7, lines 17-34, specifically lines 26-34)

Regnier does disclose comparing the bound remote key and the corresponding first address supplied by the second host processor endnode to the bound remote key and corresponding first address in the first host processor endnode (column 5, line 66-column 6, line 61, specifically column 6, lines 35-44 and 54-61)

It would have been obvious to one of ordinary skill in the art at the time the invention was made to combine the teachings

Art Unit: 2143

of these references since Regnier discloses that comparing the bound remote key and corresponding first address supplied by the second host processor endnode to the bound remote key and corresponding first address in the first host processor endnode enables verification that the information to be transferred is correct and prevents transaction of data if there is a mismatch (column 5, line 66-column 6, line 61, specifically column 6, lines 35-44 and 54-61). In view of these specific advantages and that the references are directed to performing remote direct memory accesses between endnodes over a communication fabric, one of ordinary skill would have been motivated to combine these references and would have considered them to be analogous to one another based on their related fields of endeavor, which would lead one of ordinary skill to reasonably expect a successful combination of the teachings.

Regarding claim 4, Futral and Regnier disclose the method of claim 3.

Futral does not expressly disclose wherein if the bound remote key and corresponding first address supplied by the second host processor endnode do not match the bound remote key and first address in the first host processor endnode, the second host processor endnode is not granted access to the contiguous memory address range, however, Futral does disclose

Art Unit: 2143

using the bound remote key and the corresponding first address for validation purposes (column 5, lines 21-23 and 32-42; column 7, lines 17-34, specifically lines 26-34)

Regnier discloses that if the bound remote key and corresponding first address supplied by the second host processor endnode do not match the bound remote key and first address in the first host processor endnode, the second host processor endnode is not granted access to the contiguous memory address range (column 5, line 66-column 6, line 61, specifically column 6, lines 35-44 and 54-61).

Claim 4 is rejected since the motivations regarding the obviousness of claim 3 also apply to claim 4.

Regarding claim 5, Futral and Regnier disclose the method of claim 3.

Futral does not expressly disclose wherein if the contiguous memory address range represented by the first address bound to the remote key supplied by the second host processor endnode is invalid, the second host processor endnode is not granted access to the contiguous memory address range, however, Futral does disclose using the bound remote key and the corresponding first address for validation purposes (column 5, lines 21-23 and 32-42; column 7, lines 17-34, specifically lines 26-34)

Art Unit: 2143

Regnier discloses wherein if the contiguous memory address range represented by the first address bound to the remote key supplied by the second host processor endnode is invalid, the second host processor endnode is not granted access to the contiguous memory address range (column 5, line 66-column 6, line 61, specifically column 6, lines 35-44 and 54-61).

Claim 5 is rejected since the motivations regarding the obviousness of claim 3 also apply to claim 5.

Regarding claim 8, Futral discloses the method of claim 7.

Futral does not expressly disclose the method further comprising comparing the bound remote key and the corresponding first virtual address supplied by the second host processor endnode to the bound remote key and corresponding first virtual address in the first host processor endnode and handling a page fault condition in the first host processor endnode caused by the first virtual address bound to the remote key supplied by the second host processor endnode not being previously mapped by an operating system of the first host processor endnode, however, Futral does disclose using the bound remote key and the corresponding first address for validation purposes (column 5, lines 32-42; column 7, lines 17-34, specifically lines 26-34).

Regnier discloses comparing the bound remote key and the corresponding first virtual address supplied by the second host

Art Unit: 2143

processor endnode to the bound remote key and corresponding first virtual address in the first host processor endnode and handling a page fault condition ("memory protection fault") in the first host processor endnode caused by the first virtual address bound to the remote key supplied by the second host processor endnode not being previously mapped by an operating system of the first host processor endnode (column 5, line 66-column 6, line 61, specifically column 6, lines 35-44 and 54-61).

Claim 8 is rejected since the motivations regarding the obviousness of claim 3 also apply to claim 8.

Claims 21-22 and 25 are also rejected since these claims recite a distributed computer system that contains substantially the same limitations as recited in claims 4-5 and 8 respectively.

Claims 14-15, 18, 30-31, and 34 are rejected under 35 U.S.C. 103(a) as being unpatentable over Futral et al. in view of US Patent 6 360 220 to Forin.

Regarding claim 14, Futral discloses the method of claim 1.

Futral does not expressly disclose the method further comprising unbinding the remote key from the first address with a consumer process employing an unbind remote key verb, however, Forin does disclose this limitation (column 21, line 66-column

Art Unit: 2143

22, line 19, specifically column 22, lines 16-19) (see also Figure 8)

It would have been obvious to one of ordinary skill in the art at the time the invention was made to combine the teachings of these references since Forin discloses that unbinding the remote key enables the remote key that is no longer needed to be removed to allow for other remote keys to be stored (column 23, lines 29-32). In view of these specific advantages and that the references are directed to using binded remote keys in order to transfer data over a communication fabric using a consumer process, one of ordinary skill would have been motivated to combine these references and would have considered them to be analogous to one another based on their related fields of endeavor, which would lead one of ordinary skill to reasonably expect a successful combination of the teachings.

Regarding claim 15, Futral discloses the method of claim 13.

Futral does not expressly disclose the method further comprising retiring at least one remote key that was previously obtained via the allocate remote key verb with the consumer process employing a deallocate remote key verb, however, Forin does disclose this limitation (column 21, line 66-column 22,

Art Unit: 2143

line 19, specifically column 22, lines 16-19) (see also Figure 8 and 10)

Claim 15 is rejected since the motivations regarding the obviousness of claim 14 also apply to claim 15.

Regarding claim 18, Futral discloses the method of claim 1.

Futral does not expressly disclose the method further comprising disabling a translation for the remote key after the remote key is used for the remote direct memory access operation from the second host processor endnode, however, Forin does disclose this limitation (column 21, line 66-column 22, line 19, specifically column 22, lines 16-19) (see also Figure 8)

It would have been obvious to one of ordinary skill in the art at the time the invention was made to combine the teachings of these references since Forin discloses that disabling a translation for the remote key after the remote key is used for the remote direct memory access operation from the second host processor endnode enables the remote key to be removed to allow for other remote keys to be stored (column 23, lines 29-32). In view of these specific advantages and that the references are directed to using binded remote keys in order to transfer data over a communication fabric using a consumer process, one of ordinary skill would have been motivated to combine these references and would have considered them to be analogous to one

Art Unit: 2143

another based on their related fields of endeavor, which would lead one of ordinary skill to reasonably expect a successful combination of the teachings.

Claims 30-31 and 34 are also rejected since these claims recite a distributed computer system that contain substantially the same limitations as recited in claims 14-15 and 18 respectively.

(10) Response to Argument

The Applicant argues that Futral does not teach or suggest performing a remote direct memory access operation from the second host processor endnode with a second consumer process stored in the second memory to access the contiguous memory address range including sending the bound remote key and the first address from the second host processor endnode to the first host processor endnode on the communication fabric via the second NIC and the first NIC as recited in claims 1 and 19.

The Examiner disagrees in view of the teachings of Futral. Futral clearly discloses that "The mechanism for the I/O device to transfer data beyond the I/O unit's physical memory domain is the ability to create a Remote Direct Memory Access (RDMA) object. A RDMA object identifies memory registered by a process that is accessible by a remote transport agent. The I/O device creates a RDMA object specifying the platform ID. The RDMA

Art Unit: 2143

allows the I/O device to directly use the transport for transferring data between the I/O device and a remote process.

If a suitable virtual interface (VI) does not exist, the I/O Unit creates a VI and connects it to the remote process. Once the RDMA object is created, the I/O device uses it to access the remote transport capabilities of the I/O Unit and move data directly to and from the originator's buffers." (see column 7, lines 5-17) Futral further discloses that "The SAN NICs interconnect all host systems and I/O units through the SAN Fabric. The SAN NICs provide protection so a remote unit can only access memory that it has been authorized to access. The SAN hardware can send data with no extra data copies. For a read operation, the sender tells the SAN at a local end both the source location of data in the local unit to be transferred and the destination location in a remote unit. The SAN directly moves the data across the SAN Fabric from the local unit into the desired place in the remote unit without further copies of data being made. To do this, it requires the destination unit to register the memory where the data resides with the SAN NIC. The SAN NIC returns a memory handle as a memory protection key. The SAN NIC sends the memory handle and the virtual address of the data to the remote unit to initiate the data transfer. The memory handle and virtual address are sufficient for accessing

Art Unit: 2143

memory in a particular unit. Additionally, for access to a remote unit, a platform identifier is used to identify the unit where the memory handle and virtual address are valid. Thus, the combination of a platform identifier, a memory handle for registered memory, and a virtual address uniquely identifies memory located anywhere in the clustered computer system." (see column 5, lines 18-42)

Futral also discloses that "FIG. 2 is a diagram illustrating multiple host systems accessing multiple I/O units via a SAN fabric according to an embodiment of the present invention. Cluster Host Systems 1, 2 and N, numbered 20, 22, and 24, respectively, communicate with each other via System Area Network (SAN) Fabric 26. Although three host systems are shown in FIG. 2, any number of host systems may be interconnected by the SAN Fabric. Host systems may also be thought of as servers. A server is a compute node executing an appropriate high level operating system. However, the term is generic and could also apply to other computers such as desktop personal computers (PCs) and workstations. Each Cluster Host System includes a SAN Network Interface Controller (NIC). The SAN NIC attaches to any local bus within a host or I/O unit. Host systems and I/O units are generically called units. A unit may have multiple SAN NICs installed. SAN NICs attach a unit to the SAN Fabric. The SAN NIC

Art Unit: 2143

provides connections to other units that are external to the unit." (see column 4, lines 31-50).

Futral also discloses that "The operation of the SAN NIC is abstracted by a Remote Transport Agent (RTA) (not shown). The RTA is the component in each unit that manages connections between the unit and other units and provides transport resources to exchange messages and data between various units. It abstracts reliable point-to-point communication between pairs of host units and IOPs in a SAN. Through the RTA, a host unit establishes a logical connection with an IOP with which it needs to communicate. Such connections may be created once and maintained for every possible communication endpoint, or created and deleted dynamically as needed to support I/O request operations. An IOP has exactly one logical communication endpoint to each host unit, independent of the number of SAN NICs actually configured." (see column 5, lines 51-65)

Therefore, Futral clearly discloses performing a remote direct memory access operation using an "RDMA object" from the second host processor endnode or "local unit" with a second consumer process or "RTA" stored in the second memory to access the contiguous memory address range or "memory located anywhere within the clustered computer system" including sending the bound remote key or "memory handle" and the first address or

Art Unit: 2143

"virtual address" from the second host processor endnode to the first host processor endnode or "remote unit" on the communication fabric via the second NIC and the first NIC which are disposed on the endnodes.

The Applicant also argues that a host processor endnode as defined in the specification is in "no way equivalent" to an I/O device or I/O unit.

The specification discloses:

"Distributed computer system 30 includes a system area network (SAN) 32 which is a high-bandwidth, low-latency network interconnecting nodes within distributed computer system 30. A node is herein defined to be any device attached to one or more links of a network and forming the origin and/or destination of messages within the network. In the example distributed computer system 30, nodes include host processors 34a-34d; redundant array independent disk (RAID) subsystem 33; and I/O adapters 35a and 35b. The nodes illustrated in Figure 1 are for illustrative purposes only, as SAN 32 can connect any number and any type of independent processor nodes, I/O adapter nodes, and I/O device nodes. Any one of the nodes can function as an endnode, which is herein defined to be a device that originates or finally consumes messages or frames in the distributed computer system." (page 6, lines 13-24)

Art Unit: 2143

In view of the disclosures of Futral, specifically in column 4, lines 31-50, that a host processor endnode as defined in the specification is equivalent to the "unit" or "host" which may be a computing device such as a "server" or "desktop personal computers" or "workstation". Therefore, Futral does disclose such an "endnode".

The Applicant also argues that the claims distinctly define the endnode as having a processor and memory and that they are separate and distinct from the NIC within the endnode. Futral clearly discloses such a system as both the NIC and endnode inherently contain a processor and memory to perform their respective functions.

Therefore, it is submitted that Futral does disclose these limitations of claims 1 and 19 and the rejections under Futral should be sustained. Also, the rejections under Futral in view of Forin and Regnier should be sustained.

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

Art Unit: 2143

For the above reasons, it is believed that the rejections should be sustained.

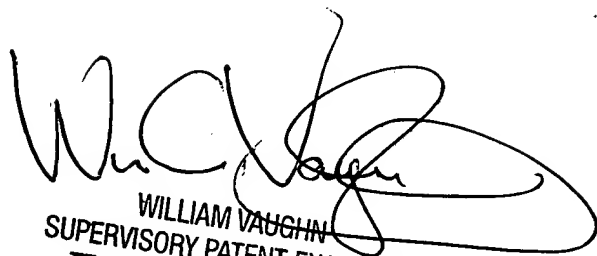
Respectfully submitted,

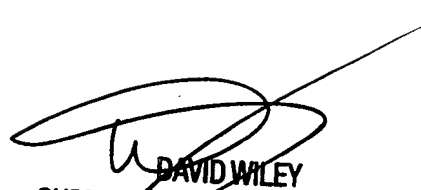
/George C. Neurauter, Jr./

George C. Neurauter, Jr.

Patent Examiner

Conferees:


WILLIAM VAUGHN
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100


DAVID WILEY
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100